SEAGATE

FIRECUDA™

Seagate® FireCuda® 540 SSD

Product Manual

Form Factor	User Capacity	Standard Models
M.2 2280-D2	1000 GB	ZP1000GM30004
	2000 GB	ZP2000GM30004

204962300, Rev A
May 2023

Contents

Support

For Internal SSD Support, visit: https://www.seagate.com/support/products/

For Firmware Download and Tools Download for Secure Erase, visit: https://www.seagate.com/support/downloads/

For information regarding online support and services, visit: http://www.seagate.com/contacts/

For information regarding Warranty Support, visit: http://www.seagate.com/support/warranty-and-replacements/

For information regarding data recovery services, visit:
http://www.seagate.com/services-software/seagate-recovery-services/recover/

For Seagate OEM and Distribution partner and Seagate reseller portal, visit: http://www.seagate.com/partners

1. Introduction

The Seagate® FireCuda® 540 SSD is a gaming NVMe SSD with PCIe Gen5 x4 interface. It is up to 17x faster than SATA SSDs and delivers ultra-fast performance and enhanced endurance for long term use.

Table 1   The FireCuda 540 SSD Features

Feature	Description	
Capacity (User)	■ 1000 GB, 2000 GB	
Certifications, Eco-Compliance	■ CE, UL, FCC, BSMI, KCC, Microsoft WHQL, VCCI, CB ■ RoHS	
Dimensions	Standard Models—1000 GB, 2000 GB ■ Length, Max 80.15 mm ■ Width, Max 22.15 mm ■ Height, Max 3.58 mm	
Form Factor	■ M2 2280-D2-M	
Weight	Standard Models ■ 1000 GB: 7.4 g ■ 2000 GB: 7.4 g	
Endurance	Total Bytes Written ■ 1000 GB: 1000 TB ■ 2000 GB: 2000 TB	See Section 2.5, *Reliability/Endurance*.
Interface Compliance	■ NVMe 2.0 ■ PCI Express Base 5.0, PCIe Gen 5 x 4 lane, and backward compatible to PCIe Gen4, Gen 3, Gen 2, and Gen 1	
NAND	■ TLC	
Operating Systems	■ Windows 10 (64 bit), Windows 11 ■ Ubuntu 16.04, 18.04 ■ CentOS 6, 7	
Performance Random	■ Read: Up to 1,300,000 IOPS ■ Write: Up to 1,500,000 IOPS	Actual performance might vary depending on use conditions and environment. See Section 2.2, *Performance*.
Performance Sequential	■ Read: Up to 10,00MB/s ■ Write: Up to 10,000MB/s	Actual performance might vary depending on the capacity, use conditions and environment. See Section 2.2, *Performance*.
Power Consumption	■ Active Power, Average: <11 W ■ Idle Power PS3, Average: <144 mW	See Section 2.3, *Power Consumption*.

Table 1   The FireCuda 540 SSD Features *(continued)*

Feature	Description	
Power Management	■ Supports ActiveStatePower Management (ASPM) ■ Supports Autonomous Power StateTransition (APST) ■ Supports L1.2	
Reliability	■ End-to-end data path protection ■ MTBF: 1.8 million hours ■ UBER: 1 error in 10^{16} bits read	
Security	■ AES-256 and TCG Opal 2.01	See Section 7.2, *Advanced Device Security Features*.
Shock and Vibration	Shock ■ Non-Operating: 1,500 G, at 0.5 ms	**See Section 2.4, *Environmental Conditions*.**
	Vibration ■ Non-Operating: 1.52 G_{RMS}, (20 to 80 Hz, Frequency)	
Temperature Range	■ Operating: 0°C to 70°C ■ Non-operating: -40°C to 85°C	
Voltage	■ Min = 3.14V±5% ■ Max = 3.47V±5%	
Warranty	Five years, or when the device reaches Host TBW, whichever happens first. Endurance rating valid for SSD Life Remaining > 1%.	

1.1   References

In case of conflict between this document and the following reference documents, this document takes precedence.

■ PCIe Specifications
   — PCIe - PCI Express Electromechanical specification, revision 4.0
   — NVMe - Non Volatile Memory Express specification 2.0
   — PCIe CEM - PCI Express Card Electromechanical specification, revision 1.1
   — PCI Express M.2 Specification, revision 1.1
■ Seagate Downloads are available on the Seagate Support page here:
   https://www.seagate.com/support/

2. Specifications

2.1 Models and Capacity

Table 2   Models and Capacity

User Capacity	Standard Models
1000 GB	ZP1000GM30004
2000 GB	ZP2000GM30004

NOTE   About capacity:

- Sector Size: 512 Bytes (default) and 4K

- User-addressable LBA count = (97696368) + (1953504 x (Desired Capacity in Gb-50.0)) From International Disk Drive Equipment and Materials Association (IDEMA) (*LBA1-03_standard.doc*)

2.2 Performance

Table 3   Random and Sequential Read and Write Performance

Capacity	CrystalDiskMark			
	Read Sequential (MB/s)	Write Sequential (MB/s)	Read Random (IOPS)	Write Random (IOPS)
1000GB	9,500	8,500	1,300,000	1,500,000
2000GB	10,000	10,000	1,490,000	1,500,000

NOTE   About performance:

- Fresh out of box (FOB) performance obtained on newly formatted drive. Performance may vary based on the SSD's firmware version, system hardware, and configuration.

- Performance based on Intel 13th Gen Core i9-13900K Platform, DRAM (5200MHz) 32GB, Windows OS.

- Performance based on CrystalDiskMark v.8.0.4

  - Sequential Read / Write: 1GB range, 1MB data size, QD=8, Thread=1
  - Random Read / Write: 1GB range, 4KB data size, QD=32, Thread=16

- For optimal performance, standard models require use of a heatsink (motherboard or discreet)

2.3   Power Consumption

Table 4   Power Consumption

Capacity	Power Consumption		
	Max Avg Read (W)	**Max Avg Write (W)**	**Idle PS3 (mW)**
1000 GB	10	10	144
2000 GB	10.5	11.5	144

NOTE   About power consumption:

- Measured using CrystalDiskMark 8.0.4, 1GB range, QD=8, Thread=1 during sequential read and sequential write on Intel 12th Gen Core i9-12900K 3.19GHz Platform + DRAM (4800MHz).

- The measured power voltage is 3.3 V.

- Measured under ambient temperature.

- Power consumption of PS3 is under LPM2+L1.0 state.

- Power Consumption can differ with flash configuration and platform.

2.4   Environmental Conditions

Table 5   Temperature, Humidity, Shock

Specification	Value

Temperature
        Operating (case temperature at specific airflow)
        Non-operating
	0℃ to 70℃
	-40℃ to 85℃

Humidity
        Operating
        Non-operating (storage)
	90%
	93%

Shock Non-operating	1,500 G, duration 0.5 ms
Vibration Non-operating	1.52 $G_{RMS,}$ (20Hz to 80Hz, Frequency)

NOTE   **Temperature** is measured without condensation. Operating mode temperature is measured by temperature sensor, SMART Attribute.

Airflow is suggested. Airflow allows the device to be operated at the appropriate temperature for each component during heavy workloads environments.

Shock and vibration results assume that the SSD is mounted securely with the input vibration applied to the SSD mounting. These specifications do not

cover connection issues that may result from testing at this level. The measured specification is in root mean square (RMS) form.

- **Non-operating Shock.** The limits of non-operating shock applies to all conditions of handling and transportation. This includes both isolated SSD and integrated SSDs. Shock may be applied in the X, Y, or Z-axis.

- **Non-Operating Vibration.** The limits of non-operating vibration shall apply to all conditions of handling and transportation. This includes both isolated SSD and integrated SSDs. Vibration may be applied in the X, Y, or Z-axis.

2.5   Reliability/Endurance

Table 6   Reliability/Endurance

Specification	Value
Mean time between failures (MTBF)	1.8 million hours
Bit Error Rate	1 error in 10^{16} bits read
Endurance	Total Bytes Written ■ 1000 GB: 1000 TB ■ 2000 GB: 2000 TB

NOTE   About endurance:

- The SSD achieves the specified MTBF in an operational environment that complies with the operational temperature range specified in this manual. Operating temperatures are measured by temperature sensor.

- Endurance rating valid for SSD Life Remaining > 1%.

3. Mechanical Dimensions and Drawings

This section includes, weight, dimensions, and mechanical drawings.

3.1 Standard Model Dimensions and Drawings

Table 7   FireCuda 540 SSD Weight and Dimensions

Capacity	Weight (g)	Length (Max. mm)	Width (Max. mm)	Height (Max. mm)
1000 GB	7.4	80.15	22.15	3.58
2000 GB	7.4	80.15	22.15	3.58

Figure 1 FireCuda 540 SSD M2 2280-D2-M Top View
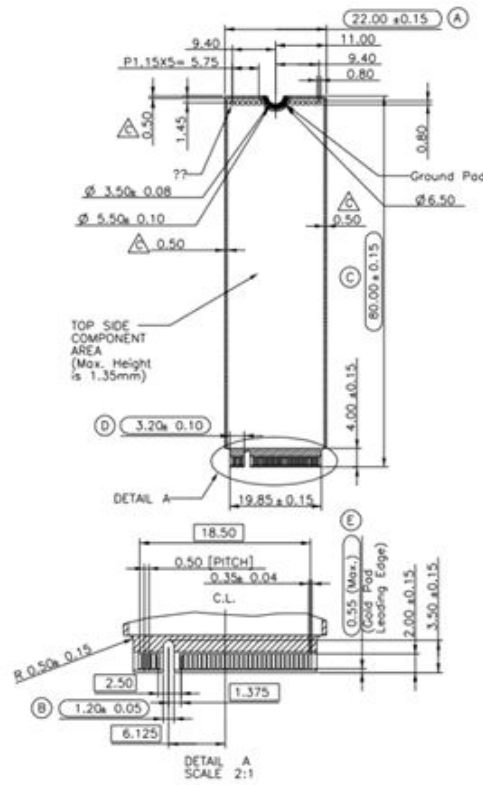
Top View

Figure 2 FireCuda 540 SSD M2 2280-D2-M Side View
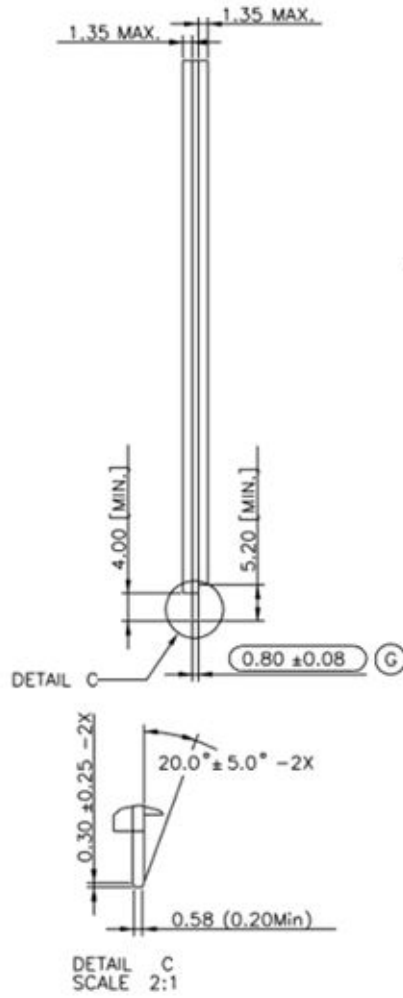
Side View

Figure 3 FireCuda 540 SSD M2 2280-D2-M Bottom View
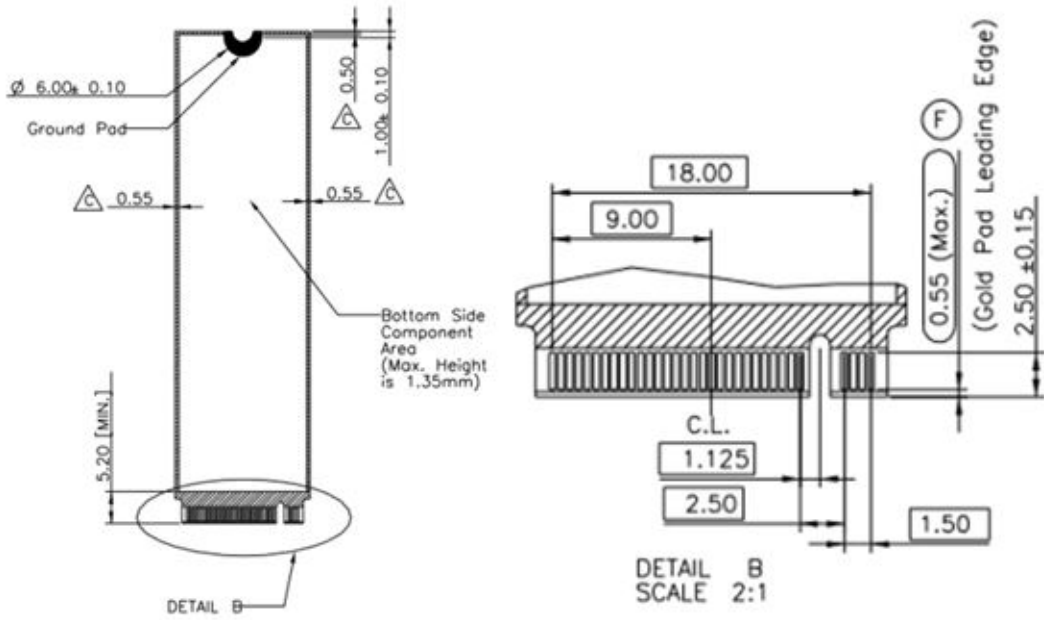
Bottom View

4. Pin and Signal Descriptions

Table 9   Pin Descriptions

Pin No.	PCIe Pin	Description
1	GND	CONFIG_3 = GND
2	3.3V	3.3V source
3	GND	Ground
4	3.3V	3.3V source
5	PETn3	PCIe TX Differential signal defined by the PCI Express M.2 spec
6	N/C	No connect
7	PETp3	PCIe TX Differential signal defined by the PCI Express M.2 spec
8	N/C	No connect
9	GND	Ground
10	LED1#	Open drain, active low signal. These signals are used to allow the add- in card to provide status indicators via LED devices that will be provided by the system.
11	PERn3	PCIe RX Differential signal defined by the PCI Express M.2 spec
12	3.3V	3.3V source
13	PERp3	PCIe RX Differential signal defined by the PCI Express M.2 spec
14	3.3V	3.3V source
15	GND	Ground
16	3.3V	3.3V source
17	PETn2	PCIe TX Differential signal defined by the PCI Express M.2 spec
18	3.3V	3.3V source
19	PETp2	PCIe TX Differential signal defined by the PCI Express M.2 spec
20	N/C	No connect
21	GND	Ground
22	N/C	No connect
23	PERn2	PCIe RX Differential signal defined by the PCI Express M.2 spec
24	N/C	No connect
25	PERp2	PCIe RX Differential signal defined by the PCI Express M.2 spec
26	N/C	No connect
27	GND	Ground
28	N/C	No connect
29	PETn1	PCIe TX Differential signal defined by the PCI Express M.2 spec
30	N/C	No connect

Table 9   Pin Descriptions *(continued)*

Pin No.	PCIe Pin	Description
31	PETp1	PCIe TX Differential signal defined by the PCI Express M.2 spec
32	N/C	No connect
33	GND	Ground
34	N/C	No connect
35	PERn1	PCIe RX Differential signal defined by the PCI Express M.2 spec
36	N/C	No connect
37	PERp1	PCIe RX Differential signal defined by the PCI Express M.2 spec
38	N/C	No connect
39	GND	Ground
40	SMB_CLK (I/O)(0/1.8V)	SMBus Clock; Open Drain with pull-up on platform
41	PETn0	PCIe TX Differential signal defined by the PCI Express M.2 spec
42	SMB_DATA (I/O)(0/1.8V)	SMBus Data; Open Drain with pull-up on platform.
43	PETp0	PCIe TX Differential signal defined by the PCI Express M.2 spec
44	ALERT#(O) (0/1.8V)	Alert notification to master; Open Drain with pull-up on platform; Active low.
45	GND	Ground
46	N/C	No connect
47	PERn0	PCIe RX Differential signal defined by the PCI Express M.2 spec
48	N/C	No connect
49	PERp0	PCIe RX Differential signal defined by the PCI Express M.2 spec
50	PERST#(I)(0/3.3V)	PE-Reset is a functional reset to the card as defined by the PCIe Mini CEM specification.
51	GND	Ground
52	CLKREQ#(I/O)(0/3.3V)	Clock Request is a reference clock request signal as defined by the PCIe Mini CEM specification; Also used by L1 PM Sub-states.
53	REFCLKn	PCIe Reference Clock signals (100 MHz) defined by the PCI Express M.2 spec.
54	PEWAKE#(I/O)(0/3.3V)	PCIe PME Wake. Open Drain with pull up on platform; Active Low.
55	REFCLKp	PCIe Reference Clock signals (100 MHz) defined by the PCI Express M.2 spec.
56	Reserved for MFG DATA	Manufacturing Data line. Used for SSD manufacturing only. Not used in normal operation. Pins should be left N/C in platform Socket.

Table 9   Pin Descriptions *(continued)*

Pin No.	PCIe Pin	Description
57	GND	Ground
58	Reserved for MFG CLOCK	Manufacturing Clock line. Used for SSD manufacturing only. Not used in normal operation. Pins should be left N/C in platform Socket.
59	Module Key M	Module Key
60	Module Key M	
61	Module Key M	
62	Module Key M	
63	Module Key M	
64	Module Key M	
65	Module Key M	
66	Module Key M	
67	N/C	No connect
68	SUSCLK(32KHz) (I)(0/3.3V)	32.768 kHz clock supply input that is provided by the platform chipset to reduce power and cost for the module.
69	N/C	PEDET (NC-PCIe)
70	3.3V	3.3V source
71	GND	Ground
72	3.3V	3.3V source
73	GND	Ground
74	3.3V	3.3V source
75	GND	Ground

5. NVMe Commands

Table 10 Admin Commands

Identifier	O/M	Command Description	Supported
00h	M	Delete I/O Submission Queue	Supported
01h	M	Create I/O Submission Queue	Supported
02h	M	Get Log Page	Supported
04h	M	Delete I/O Completion Queue	Supported
05h	M	Create I/O Completion Queue	Supported
06h	M	Identify	Supported
08h	M	Abort	Supported
09h	M	Set Feature	Supported
0Ah	M	Get Feature	Supported
0Ch	M	Asynchronous Event Request	Supported
10h	O	Firmware Commit	Supported
11h	O	Firmware Image Download	Supported
14h	O	Device Self-test	Supported
80h	O	Format NVM	Supported
81h	O	Security Send	Supported
82h	O	Security Receive	Supported
84h	O	Sanitize	Supported

Table 11 I/O Commands

Identifier	O/M	Command Description	Supported
00h	O	Flush	Supported
01h	O	Write	Supported
02h	O	Read	Supported
04h	O	Write Uncorrectable	Not Supported
05h	O	Compare	Supported
08h	O	Write Zeroes	Supported
09h	O	Dataset Management	Supported

Table 12 Set Feature Commands

Identifier	O/M	Command Description	Supported
00h		Reserved	
01h	M	Arbitration	Supported
02h	M	Power Management	Supported
03h	O	LBA Range Type	Not Supported
04h	M	Temperature Threshold	Supported
05h	M	Error Recovery	Supported
06h	O	Volatile Write Cache	Supported
07h	M	Number Of Queues	Supported
08h	M	Interrupt Coalescing	Supported
09h	M	Interrupt Vector Configuration	Supported
0Ah	M	Write Atomicity Normal	Supported
0Bh	M	Asynchronous Event Configuration	Supported
0Ch	O	Autonomous Power State Transition	Supported
0Dh	O	Host Memory Buffer	Not Supported
0Eh	O	Timestamp	Supported
10h	O	Host Controlled Thermal Management	Supported
11h	O	Non-Operational Power State Config	Supported
12h	O	Read Recovery Level Config	Not Supported
13h	O	Predictable Latency Mode Config	Not Supported
14h	O	Predictable Latency Mode Window	Not Supported
15h	O	LBA Status Information Report Interval	Not Supported
16h	O	Host Behavior Support	Not Supported
17h	O	Sanitize Config	Supported
18h	O	Endurance Group Event Configuration	Not Supported
19h	O	I/O Command Set Profile	Not Supported
1Ah	O	Spinup Control	Not Supported
1Bh to 1Fh		Reserved	
20h	O	Refer to the Key Value Command Set	Not Supported
21h to 77h		Reserved	
78h to 7Ch		Reserved	
7Dh	O	Enhanced Controller Metadata	Not Supported
7Eh	O	Controller Metadata	Not Supported
7Fh	O	Namespace Metadata	Not Supported
80h	O	Software Progress Marker	Supported
81h	O	Host Identifier	Not Supported
82h	O	Reservation Notification Mask	Not Supported
83h	O	Reservation Persistence	Not Supported
84h	O	Namespace Write Protection Config	Not Supported
85h to BFh		Reserved	

Table 13 Get Log Page Commands

Identifier	O/M	Command Description	Supported
00h	M	Supported Log Pages	Supported
01h	M	Error Information	Supported
02h	M	SMART / Health Information	Supported
03h	M	Firmware Slot Information	Supported
04h	O	Changed Namespace List	Not Supported
05h	O	Commands Supported and Effects	Supported
06h	O	Device Self-test	Supported
07h	O	Telemetry Host-Initiated	Supported
08h	O	Telemetry Controller-Initiated	Supported
09h	O	Endurance Group Information	Not Supported
0Ah	O	Predictable Latency Per NVM Set	Not Supported
0Bh	O	Predictable Latency Event Aggregate	Not Supported
0Ch	O	Ansymmetric Namespace Access	Not Supported
0Dh	O	Persistent Event Log	Supported
0Eh	O	LBA Status Information	Not Supported
0Fh	O	Endurance Group Event Aggregate	Not Supported
10h	O	Media Unit Status	Not Supported
11h	O	Supported Capacity Configuration List	Not Supported
12h	M	Feature Identifiers Supported and Effects	Supported
13h	M	NVMe-MI Commands Supported and Effects	Supported
14h	O	Command and Feature Lockdown	Not Supported
15h	O	Boot Partion	Not Supported
16h	O	Rotational Media Information	Not Supported
17h to 6Fh		Reserved	
70h	O	Discovery	Not Supported
71h to 7Fh		Reserved	
80h	O	Reservation Notification	Not Supported
81h	O	Sanitize Status	Supported
82h – BEh		Reserved	
BFh	O	Refer to the Zoned Namespace Command Set	Not Supported
C0h to FFh		Vendor specific	

6. SMART Support

The FireCuda 540 SSD supports the SMART command set.

6.1   SMART Attributes

The following table lists SMART Attributes and Descriptions.

Table 14 SMART Attributes (Log Identifier 02h)

Bytes Index	Bytes	Description
[0]	1	Critical Warning
[2:1]	2	Composite Temperature
[3]	1	Available Spare
[4]	1	Available Spare Threshold
[5]	1	Percentage Used
[31:6]	26	Reserved
[47:32]	16	Data Units Read
[63:48]	16	Data Units Written
[79:64]	16	Host Read Commands
[95:80]	16	Host Write Commands
[111:96]	16	Controller Busy Time
[127:112]	16	Power Cycles
[143:128]	16	Power On Hours
[159:144]	16	Unsafe Shutdowns
[175:160]	16	Media and Data Integrity Errors
[191:176]	16	Number of Error Information Log Entries
[195:192]	4	Warning Composite Temperature Time
[199:196]	4	Critical Composite Temperature Time
[201:200]	2	Temperature Sensor 1 (Current Temperature)
[203:202]	2	Temperature Sensor 2 (N/A)
[205:204]	2	Temperature Sensor 3 (N/A)
[207:206]	2	Temperature Sensor 4 (N/A)
[209:208]	2	Temperature Sensor 5 (N/A)
[211:210]	2	Temperature Sensor 6 (N/A)
[213:212]	2	Temperature Sensor 7 (N/A)
[215:214]	2	Temperature Sensor 8 (N/A)
[511:216]	296	Reserved

| NOTES | Here is (optional) contextual information for Log Identifier 02h: |

- ■ "Critical Warning [Byte 0]." This field indicates critical warnings for the state of the controller.

  - • Bit#0: Available spare is below threshold
  - • Bit#1: Temperature exceeded threshold or below an under temperature threshold
  - • Bit#2: Reliability is degraded due to excessive media or internal errors
  - • Bit#3: Media is placed in read only mode
  - • BIt#4: Volatile memory backup device has failed.
  - • Bit#5 - Bit#7: Reserved

- ■ "Available Spare [Byte 3]." This value (percentage) = 100* [(total reserved VB - consumed VB caused by early, later bad)/ total reserved VB]

- ■ "Percentage Used [Byte 5]." This value (percentage) = 100* (total VB erase count/ PE cycle for total VB).

Table 15 SMART Attributes (Log Identifier C0h)

Bytes Index	Bytes	Description
[7:0]	8	Device Capacity
[15:8]	8	User Capacity
[23:16]	8	NAND Read
[31:24]	8	NAND Write
[39:32]	8	NAND Erase Sector
[47:40]	8	SSD Life Remaining Percent D3
[55:48]	8	SSD Life Used Percent D3
[56]	1	WP Water Mark
[58:57]	2	Highest temperature
[62:59]	4	Flash UNC Error Count
[67:63]	5	Data E3D Error
[70:67]	4	PHY Error Count
[74:71]	4	Total Bad Block Count
[78:75]	4	Total Early Bad Block Count
[82:79]	4	Total Later Bad Block Count
[86:83]	4	Read Fail Count
[90:87]	4	Program Fail Count
[94:91]	4	Erase Failure Count
[102:95]	8	System Table Copy Count
[110:103]	8	Read Move Table Count
[114:111]	4	Data read retry count
[118:115]	4	RAID ECC retry count
[122:119]	4	RAID ECC failed count
[130:123]	8	Total Erase Count
[134:131]	4	Max Erase Count
[138:135]	4	Average Erase Count
[142:139]	4	Min Erase Count
[150:143]	8	Background read count
[154:151]	4	Host Write Uncorrectable Sector Count
[158:155]	4	PS3 Enter Success
[162:159]	4	PS4 Enter Success
[166:163]	4	Wear Leveling Count
[167]	1	Chip internal temperature
[169:168]	2	Thermal throttling
[171:170]	2	Thermal throttling time
[179:172]	8	FW Code Update Count
[511:181]	331	RSV

NOTES	For Log Identifier C0h:

- "SSD Life Remaining Percent D3 [Byte 47:40]"

  - Value (percentage) = 100 *[1 - (Average of the Flash's block erase count / NAND EOL erase count)]

- "SSD Life Used Percent [Byte 55:48]"

  - Value (percentage) = 100 *(Average of the Flash's block erase count / NAND EOL erase count)

7. Feature Details

7.1 Flash Management

7.1.1 Error Correction Code (ECC)

Flash memory cells will deteriorate with use, which might generate random bit errors in the stored data. Thus, FireCuda 540 SSD applies the fourth generation LDPC (Low Density Parity Check) of ECC algorithm, which can detect and correct errors that occur during read process, ensure data has been read correctly, as well as protect data from corruption.

7.1.2 Wear Leveling

NAND flash devices can only undergo a limited number of program/erase cycles, and in most cases, the flash media are not used evenly. If some areas get updated more frequently than others, the lifetime of the device would be reduced significantly. Thus, Wear Leveling is applied to extend the lifespan of NAND Flash by evenly distributing write and erase cycles across the media.

FireCuda 540 SSDs provide advanced Wear Leveling algorithm. This can efficiently spread out the flash usage through the whole flash media area. Moreover, by implementing both dynamic and static Wear Leveling algorithms, the life expectancy of the NAND flash is greatly improved.

7.1.3 Bad Block Management

Bad blocks are blocks that do not function properly or contain more invalid bits causing stored data to become unstable, and their reliability is not guaranteed. Blocks that are identified and marked as bad by the manufacturer are referred to as "Early Bad Blocks". Bad blocks that are developed during the lifespan of the flash are named "Later Bad Blocks". FireCuda 540 SSDs implement an efficient bad block management algorithm to detect the factory- produced bad blocks and manages bad blocks that appear with use. This practice prevents data being stored into bad blocks and further improves the data reliability.

7.1.4 TRIM

TRIM is a feature which helps improve the read/write performance and speed of solid-state drives (SSD). Unlike hard disk drives (HDD), SSDs are not able to overwrite existing data, so the available space gradually becomes smaller with each use. With the TRIM command, the operating system can inform the SSD which blocks of data are no longer in use and can be removed permanently. Thus, the SSD will perform the erase action, which prevents unused data from occupying blocks all the time.

7.1.5 SMART

SMART, an acronym for Self-Monitoring, Analysis and Reporting Technology, is an open standard that allows a hard disk drive to automatically detect its health and report potential failures. When a failure is recorded by SMART, users can choose to replace the drive to prevent unexpected outage or data loss. Moreover, SMART can inform users of impending failures while there is still time to perform proactive actions, such as copy data to another device.

7.1.6   Over Provisioning

Over Provisioning refers to the inclusion of extra NAND capacity in a SSD, which is not visible and cannot be used by users. With Over Provisioning, the performance and IOPS (Input/output Operations per Second) are improved by providing the controller additional space to manage P/E cycles, which enhances the reliability and endurance as well. Moreover, the write amplification of the SSD becomes lower when the controller writes data to the flash.

7.1.7   Firmware Upgrade

Firmware can be considered as a set of instructions on how the device communicates with the host. Firmware is upgraded when new features are added, compatibility issues are fixed, or read/write performance gets improved.

7.1.8   Thermal Throttling

The purpose of thermal throttling is to prevent any components in a SSD from over-heating during read and write operations. The device is designed with an on-board thermal sensor, controller on-die thermal sensor, and flash on-die thermal sensor, and with its accuracy, firmware can apply different levels of throttling to achieve the purpose of protection efficiently and proactively via SMART reading.

Table 16 Thermal Throttling

Item	Content
Smart reporting temperature	Flash composite temperature
Reference of temp. reading	On-board thermal sensor, Controller on-die thermal sensor, Flash on-die thermal sensor
tmt1 threshold	81℃ per Smart reported
tmt2 threshold	83℃ per Smart reported
Protect threshold	85℃ per Smart reported
Protect DRAM threshold	DRAM=88℃
Protect Controller threshold	CTL=115℃
Fatal threshold	CTL_Fatal=120℃, NAND_Fatal=90℃. DRAM_Fatal=90℃
Resume performance threshold	77℃ per Smart reported
Temperature polling frequency	Every 0.5 second
TMT1_state impact	COP0 add delay between both read and write
TMT2_state impact	COP0 add delay between both read and write

NOTE   For optimal performance:

- Provide sufficient airflow and cooling.
- Standard models d∨℄g[℄ geℒa XS ZV%fe`] (motherboard or discreet)

7.2 Advanced Device Security Features

7.2.1 NVMe format

Secure Erase is a standard NVMe format command and it writes all "0x00" to fully wipe all the data on the SSDs. When this command is issued, the SSD controller erases its storage blocks and returns the drive to its factory default settings.

7.2.2 Physical Presence SID

TCG OPAL defines Physical Presence SID (PSID) as a 32-character string. PSID reverses the SSD to its manufacturing setting when the SSD is set through TCG Opal (SED) or TCG Pyrite (non-SED). The PSID code is printed on the SSD label. PSID erases all data when reverting the SSD to manufacturing settings.

7.2.3 Manufacturer Secure ID

TCG OPAL defines the Manufacturer's Secure ID (MSID) as a 32-character string and is assigned in the manufacturing process. You cannot change this data through the host system. You can retrieve MSID electronically from the SSD across the interface. After getting the SSD, you must set a new password. If you do not set a new password, anyone who can reset the MSID can control the SSD. Such an attack on the SSD is called Denial of Service (DoS) because the rightful owner is locked out.

7.2.4 Sanitize Operation

The Sanitize feature is an alternative to the existing secure erase capabilities through the Format NVM command and makes a robust data security by ensuring the user data from the drive's media, caches and the Controller Memory Buffer are all wiped by the block erase operations, overwriting or destroying the encryption key. The following table illustrates the types of Sanitize Operations supported.

Table 17 Supported Sanitize Operations

Drive Security Type	Sanitize Operation			TCG Command Off		TCG Command Enable	
	Overwrite	Block Erase	Crypto Erase	PSID Revert Process	Instant Security Erase	PSID Revert Process	Instant Security Erase
SED - TCG Opal	Yes	Yes	No	No	No	Yes	Yes

NOTE Sanitize Overwrite command completion takes at least one hour per terabyte per pass. The number of passes is drive-selectable. The NVMe spec default is 16 passes. Contact Seagate Support for more detailed information.

The Security Erase feature erases all the data of the SED drive with the TCG-activated encrypted data structure by resetting the SSD with the PSID. This process happens relatively quickly. After the key is reset, the previously encrypted data cannot be accessed.

7.3   SSD Lifetime Management

7.3.1   Total Bytes Written (TBW)

TBW (total bytes written) is a measurement of the SSDs' expected lifespan, which represents the amount of data written to the device. To calculate the TBW of a SSD, the following equation is applied:

TBW = [(NAND Endurance) x (SSD Capacity)] / WAF

NAND Endurance: NAND endurance refers to the P/E (Program/Erase) cycle of a NAND flash. SSD Capacity: The SSD capacity is the specific capacity in total of a SSD.

WAF: Write Amplification Factor (WAF) is a numerical value representing the ratio between the amount of data that a SSD controller needs to write and the amount of data that the host's flash controller writes. A better WAF, which is near 1, guarantees better endurance and lower frequency of data written to flash memory.

7.3.2   Media Wear Indicator

Actual life indicator reported by SMART Attribute byte index [5], Percentage Used, recommends User to replace drive when reaching to 100%.

7.3.3   Read Only Mode (End of Life)

When drive is aged by cumulated program/erase cycles, media worn- out may cause increasing numbers of later bad block. When the number of available spare is less the threshold(5%, SMART attribute log ID 02h Byte4), the drive will notify Host through AER event and Critical Warning to enter Read Only Mode to prevent further data corruption. User should start to replace the drive with another one immediately.

7.4　An Adaptive Approach to Performance Tuning

7.4.1　Throughput

Based on the available space of the disk, the drive will regulate the read/write speed and manage the performance of throughput. When there still remains a lot of space, the firmware will continuously perform read/write action. There is still no need to implement garbage collection to allocate and release memory, which will accelerate the read/write processing to improve the performance. Contrarily, when the space is being used up, the drive will slow down the read/write processing, and implement garbage collection to release memory. Hence, read/write performance will become slower.

7.4.2　Predict & Fetch

Normally, when the Host tries to read data from the PCIe SSD, the PCIe SSD will only perform one read action after receiving one command. However, the drive applies Predict & Fetch to improve the read speed. When the host issues sequential read commands to the PCIe SSD, the PCIe SSD will automatically expect that the following will also be read commands. Thus, before receiving the next command, flash has already prepared the data. Accordingly, this accelerates the data processing time, and the host does not need to wait so long to receive data.

7.4.3　SLC Caching

The firmware design of the device currently adopts dynamic caching to deliver better performance for better endurance and consumer user experience. The SLC caching size is up to 1/3 of free capacity of the SSD.

8. Safety, Standards, and Compliance

Each Hard Drive and Solid State Drive ("device") has a product label that includes certifications that apply to that specific drive. The following information provides an overview of requirements that may apply to the drive.

> NOTE   *The most up to date information on Safety, Standards, and Compliance for this product is available in the Seagate HDD and SSD Regulatory Compliance and Safety document. You can find this document on the Seagate Support page here:*
> *https://www.seagate.com/support/*

8.1   Regulatory Model Numbers

The following regulatory model number represents all features and configurations in the series:

- STA033

9. FireCuda 540 SSD Installation Precautions

9.1 FireCuda 540 SSD Handling Instructions

There are a lot of components assembled on a single SSD device. Handle the drive with care especially when it has any WLCSP (Wafer Level Chip Scale Packaging) components such as PMIC, thermal sensor or load switch. WLCSP is a packaging technology widely used for making smaller footprints. However, any bumps or scratches may damage those ultrasmall parts so you must handle with gentle care. See Figure 9.2, *FireCuda 540 SSD Installation Instructions*

    CAUTION!   **DO NOT DROP SSD**

    CAUTION!   **INSTALL SSD WITH CARE**

    CAUTION!   **STORE SSD IN A PROPER PACKAGE**

9.2 FireCuda 540 SSD Installation Instructions

The FireCuda 540 SSD features the PCIe connector with M-key, which is compatible only with the M-key socket. See Use Case 1 in Figure 6, *M-Key M.2 Assembly Precautions*. As shown in Use Case 2, misuse may cause severe damage to an SSD including burn-out.
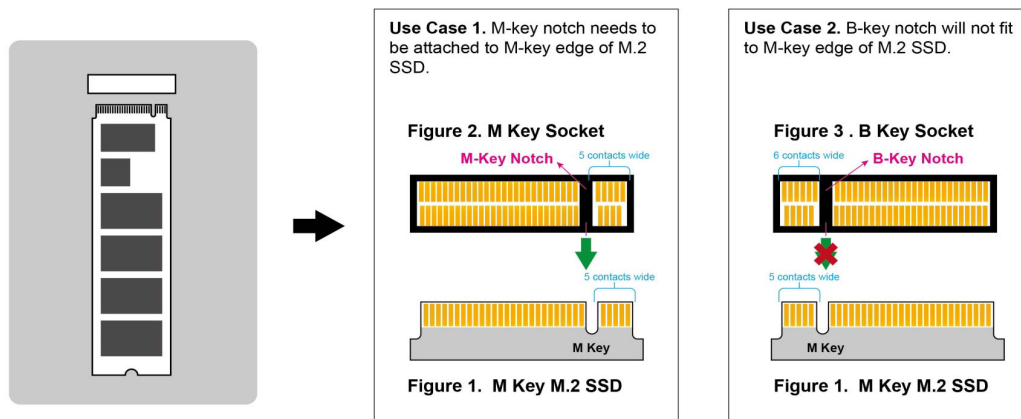

Figure 6 M-Key M.2 Assembly Precautions